SGI® UV™ Systems Linux® Configuration and Operations Guide

007–5629–005

Record of Revision

Version	Description
001	June 2010 Original publication.
002	October 2010 Updated to support the SGI Performance Suite 1.0 release.
003	February 2011 Updated to support the SGI Performance Suite 1.1 release.
004	May 2012 Updated to support the SGI Performance Suite 1.4 release.
005	November 2012 Updated to support the SGI Performance Suite 1.5 release.

Contents

Procedures

About This Guide

This guide is a reference document for people who manage the operation of SGI® UV™ 100, SGI® UV™ 1000 or SGI® UV™ 2000 systems. It explains how to perform general system configuration and operations under the Linux operating system used with SGI UV 100, SGI UV 1000, and SGI UV 2000 systems.

This manual applies to SGI UV 100, SGI UV 1000, and SGI UV 2000 systems. For the SGI UV 10 system, see the *SGI Altix UV 10 System User's Guide*. For SGI 4000 series systems, see the *Linux Configuration and Operations Guide*. For more information on SGI® ICE™ 8200 and SGI® ICE™ 8400 systems, see the *SGI Management Center for ICE*, and their respective hardware guides. For SGI® ICE X™ systems, see the *SGI Management Center for ICE X* and the *SGI ICE X System Hardware User Guide*.

SGI Management Center (SMC) software running on the system management node (SMN) provides a robust graphical interface for system configuration, operation, and monitoring. For more information on the SMC, see *SGI Management Center System Administrator Guide*.

This manual contains the following chapters:

Related Publications

For a list of manuals that support the SGI Linux releases and SGI online resources, see the *SGI Performance Suite 1.5 Start Here*.

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at: http://docs.sgi.com. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- You can view man pages by typing man *title* on a command line.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
`command`	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
variable	Italic typeface denotes variable entries and words or concepts being defined.
`user input`	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[ ]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:

  techpubs@sgi.com

- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.

- Send mail to the following address:

  SGI
Technical Publications
46600 Landing Parkway

Fremont, CA 94538

SGI values your comments and will respond to them promptly.

Configuring Your System

This chapter provides information on configuring your system and covers the following topics:

CPU Frequency Scaling

CPU frequency scaling is disabled by default on SGI UV 100 , SGI UV 1000, and SGI UV 2000 systems. This is accomplished by adding the `acpi-cpufreg` file to the `/etc/modprobe.d` directory.

For example:

```
admin:/etc/modprobe.d # cat acpi-cpufreq

# comment out the following line to enable CPU frequency scaling

install acpi-cpufreq /bin/true
```

To enable CPU frequency scaling, log into your SGI UV system (`ssh root@hostname`) and remove the `acpi-cpufreg` file in the `/etc/modprobe.d` directory. If your system is partitioned, you need to perform this on each partition.

If you decide to enable CPU frequency scaling on your system, SGI highly recommends that you set the default scaling governor to `performance` using the following script:

```
maxcpu=‘grep processor /proc/cpuinfo | awk ’{print $3}’ | tail -1‘

for cpu in ‘seq 0 $maxcpu‘

do

   cpufreq-set -c $cpu -g performance
```

```
done
```

Note: In order to enable the Intel processor's Turbo boost feature, CPU frequency scaling has to be enabled.

System Partitioning

This section describes how to partition an SGI UV 100, SGI UV 1000 or SGI UV 2000 server and contains the following topics:

Overview

A single SGI UV server can be divided into multiple distinct systems, each with its own console, root filesystem, and IP network address. Each of these software-defined group of processor cores are distinct systems and are referred to as a *partition*. Each partition can be rebooted, loaded with software, powered down, and upgraded independently. The partitions can communicate with each other over an SGI NUMAlink connection called *cross-partition communication*. Collectively, all of these partitions compose a single, shared-memory cluster.

If you enable the XPC kernel module, you enable direct memory access between partitions, which is sometimes referred to as *global shared memory*. When XPC is enabled, processes in one partition can access physical memory located on another partition. The benefits of global shared memory are currently available via SGI's Message Passing Toolkit (MPT) software. For more information on MPT, see the *Message Passing Toolkit (MPT) User Guide*.

Partition discovery software allows all of the partitions to know about each other.

Partition firewalls provide memory protection for each partition. The system software uses firewall code to open up a portion of memory so that is can be accessed by CPU cores in other partitions.

A *heartbeat* mechanism allows each partition to determine the state of all partitions in the system.

The global reference unit (GRU) no-fault code allows a partition to accesses a remote partition safely.

All of the partitions in a partitioned system have the same *system serial number*. The system serial number is stored in the system controller.

It is relatively easy to configure a large SGI UV system into partitions and reconfigure the machine for specific needs. No cable changes are needed to partition or repartition an SGI UV machine. Partitioning is accomplished by commands sent to the system controller.

Advantages of Partitioning

This section describes the advantages of partitioning an SGI UV server as follows:

- "Create a Large, Shared-memory Cluster" on page 3
- "Provides Fault Containment" on page 3
- "Allows Variable Partition Sizes" on page 4

Create a Large, Shared-memory Cluster

You can use SGI's NUMAlink technology to create a very low latency, very large, shared-memory cluster for optimized use of Message Passing Interface (MPI) software and logically shared, distributed memory access (SHMEM) routines. The globally addressable, cache coherent, shared memory is exploited by MPI and SHMEM to deliver high performance.

Provides Fault Containment

Another reason for partitioning a system is fault containment. In most cases, a single partition can be brought down (because of a hardware or software failure, or as part of a controlled shutdown) without affecting the rest of the system. Hardware memory protections prevent any unintentional accesses to physical memory on a different partition from reaching and corrupting that physical memory. For current fault containment caveats, see "Limitations of Partitioning" on page 4.

Allows Variable Partition Sizes

Partitions can be of different sizes, and a particular system can be configured in more than one way. For example, a 128-processor system could be configured into four partitions of 32 CPU cores each or configured into two partitions of 64 CPU cores each. (See "Supported Configurations" for a list of supported configurations for system partitioning.)

Your choice regarding partition sizes and the number of partitions affects both fault containment and scalability. For example, you might want to dedicate all 64 CPU cores of a system to a single large application during the night. During the day, you can partition the system into two 32–processor systems for separate and isolated use.

Limitations of Partitioning

Partitioning can increase the reliability of a system because power failures and other hardware errors can be contained within a particular partition. However, there still can be cases in which the whole shared memory cluster is affected. For example, this can occur during hardware upgrades that multiple partitions share.

If a partition is sharing its memory with other partitions, the loss of that partition may take down all other partitions that were accessing its memory. This is currently possible when an MPI or SHMEM job is running across partitions using the XPC kernel module.

Failures can usually be contained within a partition even when memory is being shared with other partitions. XPC is invoked using normal shutdown commands such as `reboot`(8) and `halt`(8) to ensure that all memory shared between partitions is revoked before the partition resets. This is also done if you remove the XPC kernel modules using the `rmmod`(8) command. Unexpected failures such as kernel panics or hardware failures almost always force the affected partition into the KDB kernel debugger or the LKCD crash dump utility. These tools also invoke XPC to revoke all memory shared between partitions before the partition resets. XPC cannot be invoked for unexpected failures such as power failures and spontaneous resets (not generated by the operating system), and thus all partitions sharing memory with the partition may also reset.

Supported SSI

The SGI UV 1000 system sizes range from 2 to 128 blades (16 to 2048 cores) in a single system image (SSI). The SGI UV 100 series is a family of multiprocessor distributed shared memory (DSM) computer systems that initially scale from 16 to 768 Intel processor cores as a cache-coherent SSI.

For SGI UV 1000 systems, the maximum number of processor cores in an SSI is 2048. The following describe the minimum and maximum metrics within an SSI:

- One partition.
- One to four racks.
- One to eight individual rack units (IRUs) with a maximum of two IRUs per rack.
- One to eight base I/O blades. Only one base I/O blade has the capability to boot the system.
- Two to 128 compute blades.
- Two to 128 SGI UV hubs. One hub resides on each compute blade.
- Two to 256 processor sockets. One socket on memory expansion blades. Two sockets on compute blades.
- 16 to 2048 processor cores, up to 4096 threads with Hyper-Threading enabled.
- Eight to 2048 DDR3 memory DIMMs (16 DIMMs maximum per compute blade)
- Up to 16 TBs, with up to 4 TB per rack when using 8 GB DIMMs.

Currently, the Linux operating system only supports 2048 cores/threads.

Note: The terms *single system image* (SSI) and *partition* can be used interchangeably.

See the *SGI Altix UV 1000 System User's Guide* for information on configurations that are supported for system partitioning.

The SGI UV 2000 system is a large, densely packed, blade-based, cache-coherent non-uniform memory access (ccNUMA), computer system that is based on the Intel® Xeon® processor E5 family. The basic building block of the UV system is the individual rack unit (IRU). The IRU is a 10U high enclosure that supports the following:

- Eight compute blades
- One chassis management controller
- Three power supplies
- Nine cooling fans

The SGI UV 2000 system scales as follows:

- From 2 to 128 compute blades in a single system image (SSI)
- A maximum of 2048 processor cores with Hyper-Threading turned off
- A maximum of 4096 processor threads (2048 processor cores) with hyper-threading turned on

  Each processor core supports two threads.

The following describe the minimum and maximum metrics within an SSI for SGI UV 2000 systems:

- The minimum granularity for a partition is two compute blades.
- Each partition must have the infrastructure to run as a standalone system. This infrastructure includes a system disk and console connection.
- An I/O blade belongs to the partition to which the attached IRU belongs. I/O blades cannot be shared by two partitions.
- Peripherals, such as dual-ported disks, can be shared the same way two nodes in a cluster can share peripherals.
- Partitions must be contiguous in the topology. For example, the route between any two nodes in the same partition must be contained within that partition and not route through any other partition. This allows intra-partition communication to be independent of other partitions.
- Partitions must be fully interconnected. That is to say, for any two partitions, there is a direct route between those partitions without passing through a third. This is required to fulfill true isolation of a hardware or software fault to the partition in which it occurs.
- If the system is unpartitioned, then routerless systems can have any number of blades, from 1-32, but the missing blades should be at the end (highest IRU highest blade number). Unpartitioned routered systems of size 3 or 4 IRUs should

be multiples of blade pairs, missing at the end. Unpartitioned routered systems of size 5 and above IRUs need to be in multiples of 4 blades, missing at the end.

See the *SGI UV 2000 System User Guide* for information on configurations that are supported for system partitioning.

For additional information about configurations that are supported for system partitioning, see your sales representative.

Installing Partitioning Software and Configuring Partitions

This section covers the following topics:

- "Enabling or Disabling Partitioning Software" on page 7
- "Partitioning a System" on page 7

Enabling or Disabling Partitioning Software

If your application uses the Message Passing Toolkit (MPT) software and uses multiple partitions, it uses kernel modules to ensure that it can access memory locations in other partitions. If you installed MPT according to the instructions in the *Message Passing Toolkit (MPT) User Guide*, the kernel modules are enabled. If the system issues the following message when your application runs, however, you need to enable the kernel modules:

```
MPT ERROR from do_cross_gets/xpmem_get, rc = -1, errno = 22
```

To enable the kernel modules, follow the installation instructions in the *Message Passing Toolkit (MPT) User Guide*.

Partitioning a System

A single SGI UV system can be divided into multiple distinct systems, each with its own console, root filesystem, and IP network address. Each of these software-defined processor groups is a distinct system referred to as a *partition*. Each partition can be rebooted, loaded with software, powered down, and upgraded independently. The partitions can communicate with each other over an SGI NUMAlink connection. Collectively, all of these partitions compose a single, shared-memory cluster. This section describes how to partition your system.

The following example shows how to use chassis manager controller (CMC) software to partition a two-rack system that contains four IRUs in four distinct systems; use the `console` command to open a console and boot each partition; and repartition it back to a single system.

Note: Each partition must have one base I/O blade and one disk blade for booting. `001i01b00` refers to rack 1, IRU 0, and blade00. `r001i01b01` refers to rack 1, IRU 0, and blade01.

The `config -v` command displays Base I/O and boot disk information. For example:

```
r001i01b00 IP93-BASEIO
r001i01b01 IP93-DISK
```

The following procedure explains how to partition your system.

Procedure 1-1 Partitioning a System Into Four Partitions

1. Use the `hwcfg` command to create four system partitions, as follows:

```
CMC:r1i1c>hwcfg partition=1 "r1i1b*"
CMC:r1i1c>hwcfg partition=2 "r1i2b*"
CMC:r1i1c>hwcfg partition=3 "r2i1b*"
CMC:r1i1c>hwcfg partition=4 "r2i2b*"
```

2. Use the `config -v` command to show the four partitions, as follows:

```
CMC:r1i1c> config -v

CMCs:          4
        r001i01c UV1000 SMN
        r001i02c UV1000
        r002i01c UV1000
        r002i02c UV1000

BMCs:         64
        r001i01b00 IP93-BASEIO P001
        r001i01b01 IP93-DISK P001
        r001i01b02 IP93-INTPCIE P001
        r001i01b03 IP93 P001
        r001i01b04 IP93 P001
        r001i01b05 IP93 P001
        r001i01b06 IP93 P001
        r001i01b07 IP93 P001
        r001i01b08 IP93 P001
```

```
r001i01b09 IP93-INTPCIE P001
r001i01b10 IP93-INTPCIE P001
r001i01b11 IP93-INTPCIE P001
r001i01b12 IP93-INTPCIE P001
r001i01b13 IP93 P001
r001i01b14 IP93 P001
r001i01b15 IP93 P001
r001i02b00 IP93-BASEIO P002
r001i02b01 IP93-DISK P002
r001i02b02 IP93-INTPCIE P002
r001i02b03 IP93 P002
r001i02b04 IP93 P002
r001i02b05 IP93 P002
r001i02b06 IP93 P002
r001i02b07 IP93 P002
r001i02b08 IP93 P002
r001i02b09 IP93 P002
r001i02b10 IP93 P002
r001i02b11 IP93 P002
r001i02b12 IP93 P002
r001i02b13 IP93 P002
r001i02b14 IP93 P002
r001i02b15 IP93 P002
r002i01b00 IP93-BASEIO P003
r002i01b01 IP93-DISK P003
r002i01b02 IP93 P003
r002i01b03 IP93 P003
r002i01b04 IP93 P003
r002i01b05 IP93 P003
r002i01b06 IP93 P003
r002i01b07 IP93 P003
r002i01b08 IP93 P003
r002i01b09 IP93 P003
r002i01b10 IP93 P003
r002i01b11 IP93 P003
r002i01b12 IP93 P003
r002i01b13 IP93 P003
r002i01b14 IP93 P003
r002i01b15 IP93 P003
r002i02b00 IP93-BASEIO P004
r002i02b01 IP93-DISK P004
```

```
        r002i02b02 IP93 P004
        r002i02b03 IP93 P004
        r002i02b04 IP93 P004
        r002i02b05 IP93 P004
        r002i02b06 IP93 P004
        r002i02b07 IP93 P004
        r002i02b08 IP93 P004
        r002i02b09 IP93 P004
        r002i02b10 IP93 P004
        r002i02b11 IP93 P004
        r002i02b12 IP93 P004
        r002i02b13 IP93 P004
        r002i02b14 IP93 P004
        r002i02b15 IP93 P004

Partitions:      4
        partition001 BMCs:   16
        partition002 BMCs:   16
        partition003 BMCs:   16
        partition004 BMCs:   16
```

3. Use the `hwcfg` command to display the four partitions, as follows:

```
CMC:r1i1c> hwcfg
NL5_RATE=5.0
PARTITION=1 ............................................... 16/64 BMC(s)
PARTITION=2 ............................................... 16/64 BMC(s)
PARTITION=3 ............................................... 16/64 BMC(s)
PARTITION=4 ............................................... 16/64 BMC(s)
```

4. Use the following command to reset the system and boot the four partitions:

   • If the power is currently off:

   ```
   CMC:r1i1c> power on "p*"
   ```

   In the preceding command, the quotation marks are required in order to prevent shell expansion.

Note: If all four partitions are to be powered on at once, you must either use the command above or else use the `--override` option. The `power on` command alone (without options or arguments) would not succeed in this instance because it would attempt to power on across partition boundaries.

- If the power is already on:

  ```
  CMC:r1i1c> power reset "p*"
  ```

5. Use the `console` command to open consoles to each partition and boot the partitions.

   The following command opens a console to partition one:

```
CMC:r1i1c> console p1
console: attempting connection to localhost...
console: connection to SMN/CMC (localhost) established.
console: requesting baseio console access at partition 1 (r001i01b00)...
console: tty mode enabled, use 'CTRL-]' 'q' to exit
console: console access established (OWNER)
console: CMC <--> BASEIO connection active
************************************************
*******  START OF CACHED CONSOLE OUTPUT  *******
************************************************

******** [20100513.215944] BMC r001i01b15: Cold Reset via NL broadcast reset
******** [20100513.215944] BMC r001i01b07: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b13: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b05: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b06: Cold Reset via NL broadcast reset
******** [20100513.215946] BMC r001i01b10: Cold Reset via NL broadcast reset
******** [20100513.215946] BMC r001i01b09: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b11: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b12: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b04: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b08: Cold Reset via NL broadcast reset
******** [20100513.215946] BMC r001i01b02: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b00: Cold Reset via NL broadcast reset
******** [20100513.215945] BMC r001i01b14: Cold Reset via NL broadcast reset
******** [20100513.215947] BMC r001i01b09: Cold Reset via ICH
******** [20100513.215946] BMC r001i01b12: Cold Reset via ICH
```

```
******** [20100513.215947] BMC r001i01b10: Cold Reset via ICH
******** [20100513.215947] BMC r001i01b11: Cold Reset via ICH
******** [20100513.215947] BMC r001i01b02: Cold Reset via ICH
******** [20100513.215947] BMC r001i01b00: Cold Reset via ICH
******** [20100513.215953] BMC r001i01b03: Cold Reset via NL broadcast reset
******** [20100513.220011] BMC r001i01b01: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b08: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b07: Cold Reset via NL broadcast reset
******** [20100513.220011] BMC r001i01b15: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b06: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b05: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b14: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b13: Cold Reset via NL broadcast reset
******** [20100513.220011] BMC r001i01b04: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b03: Cold Reset via NL broadcast reset
******** [20100513.220013] BMC r001i01b09: Cold Reset via NL broadcast reset
******** [20100513.220013] BMC r001i01b10: Cold Reset via NL broadcast reset
******** [20100513.220013] BMC r001i01b11: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b12: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b02: Cold Reset via NL broadcast reset
******** [20100513.220012] BMC r001i01b00: Cold Reset via NL broadcast reset
******** [20100513.220014] BMC r001i01b09: Cold Reset via ICH
******** [20100513.220014] BMC r001i01b10: Cold Reset via ICH
******** [20100513.220014] BMC r001i01b11: Cold Reset via ICH
******** [20100513.220013] BMC r001i01b12: Cold Reset via ICH
******** [20100513.220013] BMC r001i01b02: Cold Reset via ICH
******** [20100513.220016] BMC r001i01b00: Cold Reset via ICH
******** [20100513.220035] BMC r001i01b14: Cold Reset via NL broadcast reset
******** [20100513.220035] BMC r001i01b06: Cold Reset via NL broadcast reset
******** [20100513.220034] BMC r001i01b15: Cold Reset via NL broadcast reset
******** [20100513.220035] BMC r001i01b05: Cold Reset via NL broadcast reset
******** [20100513.220034] BMC r001i01b01: Cold Reset via NL broadcast reset
******** [20100513.220035] BMC r001i01b07: Cold Reset via NL broadcast reset
 .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  .  ....
Hit [Space] for Boot Menu.
ELILO boot:
```

6. Use the `console` command to open consoles on the other three partitions and boot them. The system then has four single system images.

7. Use the `hwcfg -c` partition command to clear the four partitions, as follows:

```
CMC:r1i1c> hwcfg -c partition
PARTITION=0
PARTITION=0
```

Note: This command can take several minutes to complete on large systems.

---

8. To reset the system and boot it as a single system image (one partition), use the following command:

```
CMC:r1i1c> power reset "p*"
```

System Operation

This chapter describes some basic system operations that use the chassis manager controller (CMC) commands. It covers the following topics:

SGI UV System Control Network

SGI Management Center (SMC) software running on the system management node (SMN) provides a robust graphical interface for system configuration, operation, and monitoring. For more information, see the *SGI Management Center System Administrator's Guide*.

The SGI UV 2000, SGI UV 1000, and the SGI UV 100 system control networks are private, closed networks. It is not to be reconfigured in any way different from the standard UV installation, nor is it to be directly connected to any other network. The SGI UV system control network does not accommodate additional network traffic, routing, address naming other than its own schema, and DCHP controls other than its own configuration. The system control network is not security hardened nor is it tolerant of heavy network traffic. It is vulnerable to Denial of Service attacks. The system management node acts as a gateway between the UV system control network and any other networks. For more information on the SGI UV system control network, see the *SGI Altix UV 100 System User's Guide*, the *SGI Altix UV 1000 System User's Guide*, or the *SGI UV 2000 System User Guide*.

Connecting to the UV System Controller Network Using the CMC

The SGI UV software controller is designed to manage and monitor the individual blades in SGI UV systems. Depending on your system configuration, you can monitor and operate the system from the system management node (SMN) or on smaller systems, such as, the SGI UV 100 from the CMC itself. SGI UV 1000 systems of up to 16 racks (four building blocks, also called one super block) can also be controlled and monitored from a CMC in the system. For more information, see the *SGI UV CMC Software User Guide*.

Power on and Booting an SGI UV System from Complete Power Off

To boot an SGI UV system from complete power off, perform the following steps:

1. Make sure the power breakers are on.
2. Establish a serial connection to the CONSOLE on the CMC or establish a network connection to the CMC. For a network connection, your PC or workstation must be connected to the CMC (via the network connection) and have its `/etc/hosts` file set up to include the CMCs. See "Connecting to the UV System Controller Network Using the CMC" on page 16.
3. Type the following command:

```
ssh root@hostname-cmc
SGI Chassis Manager Controller, Firmware Rev. 0.0.22

CMC:r1i1c>
```

   Typically, the default password set out of the factory is **`root`**. The CMC prompt appears. **CMC:r1i1c** refers to rack 1, IRU 1, CMC.

   If the host name is not set up in the PC or workstation's hosts file, you can simply use the IP address of the CMC, as follows:

```
ssh root@<IP-ADDRESS>
```

4. Power up your SGI UV system, using the `power on` command, as follows:

```
CMC:r1i1c> power on
```

5. Open a second console to the CMC, using the `console` command, to see the system power on, as follows:

```
ssh root@hostname-cmc
SGI Chassis Manager Controller, Firmware Rev. 0.0.22

CMC:r1i1c> console
console: attempting connection to localhost...
console: connection to SMN/CMC (localhost) established.
console: requesting baseio console access at r001i01b00...
console: tty mode enabled, use 'CTRL-]' 'q' to exit
console: console access established
console: CMC <--> BASEIO connection active
************************************************
*******  START OF CACHED CONSOLE OUTPUT  *******
************************************************

******** [20100512.143541] BMC r001i01b10: Cold Reset via NL broadcast reset
******** [20100512.143541] BMC r001i01b07: Cold Reset via NL broadcast reset
******** [20100512.143540] BMC r001i01b08: Cold Reset via NL broadcast reset
******** [20100512.143540] BMC r001i01b12: Cold Reset via NL broadcast reset
******** [20100512.143541] BMC r001i01b14: Cold Reset via NL broadcast reset
******** [20100512.143541] BMC r001i01b04: Cold Reset via NL
```

Note: Use `CTRL-] q` to exit the console.

6. Depending upon the size of your system, in can take 5 to 10 minutes for the SGI UV system to power on. When the **shell>** prompt appears, enter **`fs0:`**, as follows:

```
shell> fs0:
```

7. At the **fs0:** prompt, enter **`boot`**, as follows:

```
fs0:> boot
```

The ELILO Linux boot loader is called, various SGI configuration scripts are run, and the SUSE Linux Enterprise Server or Red Hat Enterprise Linux installation program appears.

Power off an SGI UV System

To power down the SGI UV sytem, use the `power off` command, as follows:

```
CMC:r1i1c> power off
==== r001i01c (PRI) ====
```

You can use the `power status` command, to check the power status of your system, as follows:

```
CMC:r1i1c> power status
==== r001i01c (PRI) ====
on: 0, off: 32, unknown: 0, disabled: 0
```

Power NMI to Drop into KDB

To send a nonmaskable interrupt (NMI) signal from the power command to the CMC to drop into the kernel debugger (KDB), use the `power nmi` command, as follows:

```
CMC:r1i1c> power nmi

Entering kdb (current=0xffff8aa3fe11c040, pid 0) on processor 7 \
due to NonMaskable Interrupt @ 0xffffffff8100ad42
     r15 = 0x0000000000000000      r14 = 0x0000000000000000
     r13 = 0x0000000000000000      r12 = 0x0000000000000000
      bp = 0xffffffff81927380       bx = 0xffff8ac1ff11dfd8
     r11 = 0xffffffff8101a2c0      r10 = 0xffff88000beefd18
      r9 = 0x00000000ffffffff       r8 = 0x0000000000000000
      ax = 0x0000000000000000       cx = 0x0000000000000000
      dx = 0x0000000000000000       si = 0xffff8ac1ff11dfd8
      di = 0xffffffff81a2b308  orig_ax = 0xffffffffffffffff
      ip = 0xffffffff8100ad42       cs = 0x0000000000000010
   flags = 0x0000000000000246       sp = 0xffff88000bee7ff0
      ss = 0x0000000000000018 &regs = 0xffff88000bee7f58
[7]kdb>
```

Upgrading System BIOS

To upgrade the compute blade PROM, perform the following steps:

1. From the CMC prompt, to show the current PROM level, type the following command:

```
CMC:r1i1c> showbios
Flashed on Sat May  1 14:14:45 UTC 2010 was bios.latest.fd (20100429_1603)
```

2. Get the newest PROM image from SupportFolio Online at http://support.sgi.com/.

3. Copy the latest BIOS to a directory on the CMC in `/work/bmc/common/`. An example directory is as follows:

```
CMC:r1i1c> ls
bios.latest.fd flashbios
```

4. Use the `flashbios` command to flash the compute blade BIOS.

```
CMC:r1i1c> flashbios
Using default bios: bios.latest.fd
Checking processor status on all nodes....
Done. System is read for BIOS flash update
Flashing bios bios.lastest.fd (20100429_1603) This will take several minutes.
...
```

For more information on using CMC software, see the *SGI UV CMC Software User Guide*.

For information on updating firmware on your SGI UV system from the system management node (SMN) and new firmware update commands, see the *SGI UV System Management Node Administrator Guide*.

Index